

# APPARATUS FOR MANAGING VIRTUALIZED- INFORMATION

## BACKGROUND OF THE INVENTION

### 1) Field of the Invention

5        The present invention relates to a virtualized-information management apparatus that manages corresponding information between real storage areas of plural storages connected to a network and a virtual storage area created by virtually integrating the real storage areas, and that efficiently builds a virtualizing storage system  
10      with reliability.

### 2) Description of the Related Art

      A technique of a storage area network (SAN) configured by connecting a plurality of storages such as a disk array drive and a tape drive via a network provides storage with large capacity to users.  
15      Recently, a virtualizing technique, where real volumes of the storages connected with the SAN are utilized as one virtual volume, has been applied to the SAN technique.

      In such a virtualizing technique, mapping information between  
20      sectors of the real volumes of the storages and sectors of the virtual volume created virtually as one storage device is created, and the input/output (I/O) process is executed based on the mapping information.

      The mapping information is stored in a management server  
25      connected with the SAN, and centrally managed by the management

server. Therefore, when a malfunctioning occurs in the management server, access to data stored in the real volumes is impossible.

Japanese Patent Application Laid-Open No. 2002-91706 discloses a technique in which a plurality of mapping agents keep the 5 mapping information in volatile memories, and a controller that manages the mapping information updates and backs up the mapping information held by the mapping agents.

In this system, even if a certain mapping agent malfunctions, the controller transmits correct mapping information to the mapping agent 10 to recover it to a normal state.

However, since the controller centrally manages the mapping information held in the mapping agents, when any abnormality occurs in the controller itself, the mapping information cannot be properly managed.

15 Particularly, when an IP-SAN, which is a network where a plurality of SANs are connected via an Internet protocol (IP) network, is built, if a failure occurs in the controller, this affects management of all the SANs.

Even if the controller has no failure, when a failure occurs in the 20 network that connects the controller, transmission/reception of data between the controller and the mapping agents cannot be executed. For this reason, the mapping information in all the mapping agents cannot be managed.

25 SUMMARY OF THE INVENTION

It is an object of the present invention to solve at least the problems in the conventional technology.

The virtualized-information management apparatus for managing corresponding information between real storage areas of a plurality of storages connected to a network and a virtual storage area built by virtually integrating the real storage areas, according to one aspect of the present invention includes a transmission/reception unit that transmits/receives, between the virtualized-information management apparatus and another virtualized-information management apparatus connected to another network, information on updating the corresponding information, and an updating unit that updates the corresponding information based on the information received.

The method of managing corresponding information between real storage areas of a plurality of storages connected to a network and a virtual storage area built by virtually integrating the real storage areas in a virtualized-information management apparatus connected to the network, according to another aspect of the present invention includes transmitting information on updating the corresponding information to another virtualized-information management apparatus connected to another network, receiving the information on updating the corresponding information from the another virtualized-information management apparatus, and updating the corresponding information based on the information received.

The computer program for managing corresponding information

between real storage areas of a plurality of storages connected to a network and a virtual storage area built by virtually integrating the real storage areas in a virtualized-information management apparatus connected to the network, according to still another aspect of the

5 present invention makes a computer executes the method according to the present invention.

The other objects, features, and advantages of the present invention are specifically set forth in or will become apparent from the following detailed descriptions of the invention when read in conjunction  
10 with the accompanying drawings.

#### BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a schematic diagram of a virtualizing storage system according to an embodiment of the present invention;

15 Fig. 2 is a block diagram of a virtualized-information management apparatus 10 shown in Fig. 1;

Fig. 3A and Fig. 3B are tables of an example of a data configuration of real volume information 230 shown in Fig. 2;

20 Fig. 4A and Fig. 4B are tables of an example of a data configuration of mapping information 231 shown in Fig. 2;

Fig. 5 is a flowchart of a mapping-information updating procedure by a mapping information management unit 240;

Fig. 6 is a flowchart of a data acquisition procedure by a data storage/acquisition unit 242;

25 Fig. 7 is a flowchart of a virtual-volume building procedure by a

virtual volume building unit 244; and

Fig. 8 is a schematic diagram of a virtualizing storage system when fiber channel internet protocol (FCIP) gateways 16a and 16b execute an encapsulating process.

5

#### DETAILED DESCRIPTION

Exemplary embodiment of a method of and an apparatus for managing virtualized-information according to the present invention are explained in detail with reference to the accompanying drawings. The 10 embodiment explains a fiber channel-storage area network (FC-SAN) in which a storage on each SAN transmits/receives data according to an FC protocol.

Fig. 1 is a schematic diagram of a virtualizing storage system according to an embodiment of the present invention.

15 The SAN 14a and the SAN 14b are connected via an IP network 15. A virtualized-information management apparatus 10a intervenes between the SAN 14a and the IP network 15, and a virtualized-information management apparatus 10b intervenes between the SAN 14b and the IP network 15.

20 The SAN 14a is configured by connecting a plurality of FC switches, not shown, with the virtualized-information management apparatus 10a, real volumes 11a and 11b, and application servers 13a and 13b. The real volumes 11a and 11b are virtualized so as to configure one virtual volume 12a.

25 Similarly, the SAN 14b is configured by connecting a plurality of

FC switches, not shown, with the virtualized-information management apparatus 10b, real volumes 11c and 11d, and application servers 13c and 13d. The real volumes 11c and 11d are virtualized so as to configure one virtual volume 12b.

5 The virtualized-information management apparatuses 10a and 10b manage mapping information (corresponding information) between sectors of the real volumes 11a to 11d and sectors of the virtual volumes 12a and 12b. The mapping information includes a relationship between the sectors of the real volumes (real storage 10 areas) and the sectors of the virtual volumes (virtual storage areas), and information for specifying virtualized-information management apparatuses that manage the virtual volumes.

The virtualized-information management apparatus 10a has mapping information 231a, and the mapping information 231a holds not 15 only the mapping information about the virtual volume 12a belonging to the SAN 14a connected with the virtualized-information management apparatus 10a but also the mapping information about the virtual volume 12b belonging to another SAN 14b.

The virtualized-information management apparatus 10b has 20 mapping information 231b, and the mapping information 231b holds not only the mapping information about the virtual volume 12b belonging to the SAN 14b connected with the virtualized-information management apparatus 10b but also the mapping information about the virtual volume 12a belonging to another SAN 14a.

25 When the application server 13a or 13b accesses to the virtual

volume 12b, it makes an access request to the virtualized-information management apparatus 10a. The virtualized-information management apparatus 10a refers to the mapping information 231a, and checks to find on which virtualized-information management apparatus on the

5 SAN the virtual volume 12b is managed.

When the virtualized-information management apparatus 10b that manages the virtual volume 12b is specified, the virtualized-information management apparatus 10a transmits the access request to the specified virtualized-information management apparatus 10b. The virtualized-information management apparatus 10b which receives the access request accesses to the virtual volume 12b managed by the apparatus 10b, and transmits the result to the application server 13a or 13b via the virtualized-information management apparatus 10a.

15 Similarly, when the application server 13c or 13d accesses to the virtual volume 12a, an access request is sent to the virtualized-information management apparatus 10b. The virtualized-information management apparatus 10b refers to the mapping information 231b, and checks to find on which

20 virtualized-information management apparatus on the SAN the virtual volume 12a is managed.

When the virtualized-information management apparatus 10a that manages the virtual volume 12a is specified, the virtualized-information management apparatus 10b transmits the

25 access request to the specified virtualized-information management

apparatus 10a. The virtualized-information management apparatus 10a which receives the access request accesses to the virtual volume 12a managed by the apparatus 10a, and transmits the result to the application server 13c or 13d via the virtualized-information

5 management apparatus 10b.

When a new virtual volume is created on the SAN 14b managed by the virtualized-information management apparatus 10b, the virtualized-information management apparatus 10a receives updated information relating to the mapping information 231b of the

10 virtualized-information management apparatus 10b, and updates its mapping information 231a based on the mapping information 231b.

Similarly, when a new virtual volume is created on the SAN 14a managed by the virtualized-information management apparatus 10a, the virtualized-information management apparatus 10b receives updated

15 information relating to the mapping information 231a of the virtualized-information management apparatus 10a, and updates its mapping information 231b based on the mapping information 231a.

The mapping information 231a of the virtualized-information management apparatus 10a and the mapping information 231b of the

20 virtualized-information management apparatus 10b can have common mapping information. Even when a new virtual volume is created on another SAN, its information can be reflected on the mapping information 231a or 231b, so that the newly created virtual volume can be utilized.

25 In this case, even when abnormality occurs in the

virtualized-information management apparatus 10a and thus it cannot refer to the mapping information 231a, the application servers 13c and 13d refer to the mapping information 231b, so as to be able to access to the virtual volume 12b regardless of the abnormality of the

5 virtualized-information management apparatus 10a.

On the contrary, when abnormality occurs in the virtualized-information management apparatus 10b and thus it cannot refer to the mapping information 231b, the application servers 13a and 13b refer to the mapping information 231a, so as to be able to access

10 to the virtual volume 12a regardless of the abnormality of the virtualized-information management apparatus 10b.

Differently from the conventional art in which the mapping information is integrally managed, the virtualized-information management apparatuses 10a and 10b hold the common mapping

15 information 231a and 231b in a dispersed manner so that the virtualizing storage system having high reliability can be configured in an FC-SAN where the two SANs 14a and 14b are connected by the IP network 15.

The case where two SANs 14a and 14b are connected by the IP 20 network 15 is explained, but its number is not limited to two and two or more SANs may be connected by the IP network 15.

Fig. 2 is a block diagram of a virtualized-information management apparatus 10 shown in Fig. 1. The virtualized-information management apparatus 10 is connected with the 25 SAN 14 and the IP network 15.

The virtualized-information management apparatus 10 according to the present invention refers to the mapping information (corresponding information), so as to execute a mapping process or the like for specifying a sector of a real volume based on a sector of the 5 specified virtual volume when the application servers 13a to 13d connected with the SANs 14a and 14b write or read data into/from the virtual volume 12a or 12b.

The virtualized-information management apparatus 10 transmits/receives updated information of the mapping information 231 10 to/from another virtualized-information management apparatus 10 connected with another SAN 14, so as to hold the common mapping information 231.

The virtualized-information management apparatus 10 has an interface unit 20, an input unit 21, a display unit 22, a storage unit 23 15 and a control unit 24. The interface unit 20 is a network interface that transmits/receives data to/from the application servers 13a to 13d and the virtualized-information management apparatus 10 connected with another SAN 14 via the SAN 14 or the IP network 15.

The input unit 21 is an input device such as a keyboard and a 20 mouse, and the display unit 22 is a display device such as a display. The storage unit 23 is storage such as a hard disk device. The storage unit 23 has real volume information 230, mapping information 231, and cache data 232.

The real volume information 230 includes information about real 25 volumes present in the SAN 14 connected with the

virtualized-information management apparatus 10. Fig. 3A and Fig. 3B are tables of an example of a data configuration of real volume information 230 shown in Fig. 2. The real volume information 230 includes real volume identifying information illustrated in Fig. 3A, and 5 real volume free capacity information illustrated in Fig. 3B.

The real volume identifying information illustrated in Fig. 3A has items for name of virtual volume, FC address of storage, logical unit number (LUN), and real volume storage capacity. The name of the virtual volume includes identifying names given to the virtual volumes 10 on the SANs 14 managed respectively by the virtualized-information management apparatus 10.

The FC address of the storage includes FC addresses of the storages having the real volumes configuring the virtual volume. The FC addresses are managed by FC switches, and the respective 15 storages can be integrally identified by the FC addresses. The LUN includes logical unit numbers for identifying the real volumes. Each virtual volume includes a plurality of logical volumes. The real volume storage capacity is a storage capacity with which the real volumes can store data therein.

20 The real volume free capacity information illustrated in Fig. 3B has items for FC address of storage, type of real volume, and real volume free capacity. The FC address of real volume corresponds to the FC address of real volume illustrated in Fig. 3A.

The type of real volume includes types of the storages 25 configuring the real volumes, and includes proprieties of the storages

such as a high-speed hard disk drive (HDD), a low-speed HDD, and a tape drive. The real volume free capacity is information including free capacities of the real volumes. When the FC address of storage is "021a00", types of two real volumes "high-speed HDD" and "low-speed 5 HDD" are related with this device. This represents that this device is configured by two kinds of HDDs.

Referring back to Fig. 2, the mapping information 231 includes the mapping information between the logical sectors of the virtual volumes and the real sectors of the real volumes and the information 10 about the virtualized-information management apparatus 10 that manages the virtual volumes. The updated information of the mapping information 231 is transmitted/received to/from the virtualized-information management apparatus 10 connected with another SAN 14, so that the virtualized-information management 15 apparatuses 10 hold the common mapping information 231.

Fig. 4A and Fig. 4B are tables of an example of a data configuration of a mapping information 231 shown in Fig. 2. The mapping information 231 has virtual volume identifying information illustrated in Fig. 4A and virtual volume sector information illustrated in 20 Fig. 4B.

The virtual volume identifying information illustrated in Fig. 4A has items for name of common virtual volume, name of virtualized-information management apparatus, IP addresses of virtualized-information management apparatus, name of virtual volume 25 in virtualized-information management apparatus, and a free capacity of

virtual volume.

The name of common virtual volume includes identifying names which are commonly used and can integrally identify the virtual volumes in any virtualized-information management apparatus 10 connected with 5 any SAN 14. The name of virtualized-information management apparatus includes identifying names for identifying the virtualized-information management apparatus 10 which is connected with the SAN 14 where the virtual volume is present and manages the mapping information 231. The IP address of virtualized-information 10 management apparatus includes IP addresses of the virtualized-information management apparatuses 10 identified by the names of the virtualized-information management apparatuses.

The name of virtual volume in virtualized-information management apparatus includes identifying names for the virtual 15 volumes corresponding to the names of the virtual volumes illustrated in Fig. 3A. These identifying names are available only for the virtualized-information management apparatus 10 which stores the real volume information 230 illustrated in Figs. 3A and 3B. The virtual volume free capacity is information which includes the free capacity of 20 the virtual volumes identified by the names of the common virtual volume.

In the example of Fig. 4A, two names of virtualized-information management apparatuses, two IP addresses of the virtualized-information management apparatuses, two names of the 25 virtual volumes in the virtualized-information management apparatuses,

and two free capacities of the virtual volume, are present in the virtual volume where the name of the common virtual volume is volume 1.

This is because in order to balance a load on the real volume, the data stored in a real volume corresponding to "volume A" are copied 5 in a real volume corresponding to "volume B". The data are copied between the volumes in such a manner, so that an access destination can be switched according to the load of the access on a volume.

The virtual volume sector information illustrated in Fig. 4B has items for name of common virtual volume, virtual volume sector, name 10 of virtual volume in virtualized-information management apparatus, FC address of storage, and real volume sector. The name of common virtual volume corresponds to the name of common virtual volume illustrated in Fig. 3A, and the virtual volume sector is a sector of a virtual volume identified by the name of the common virtual volume.

15 The name of virtual volume in the virtualized-information management apparatus corresponds to the name of the virtual volume in the virtualized-information management apparatus illustrated in Fig. 4A. The FC address of the storage corresponds to the FC address of storage illustrated in Fig. 3A. The real volume sector is a sector of the 20 real volume mapped in the virtual volume sector.

For the name of the common virtual volume and the virtual volume sector, two names of the virtual volumes in the virtualized-information management apparatuses, two FC addresses of the storages, and two real volume sectors are present. This is 25 because data stored in the real volume corresponding to "volume A" are

copied into the real volume corresponding to "volume B" to balance the load on the real volume.

When data are read, specifying of the name of the common virtual volume into which the data are stored and the virtual volume 5 sector is accepted. "Volume 1" is specified as the name of the common virtual volume, and "1 to 1,500" is specified as the virtual volume sector here.

In this case, the virtual volume sector information in Fig. 4B is referred to, and the name of the virtual volume in the 10 virtualized-information management apparatus has two kinds of names "volume A" and "volume B". When the load of the access to the real volumes configuring "volume A" is large, "volume B" is a virtual volume which is accessed instead.

When "volume 1" and "1 to 1,500" are specified as the access 15 destination of data, "volume A" as a copy source is set as an access destination. The name of the virtualized-information management apparatus that manages "volume A" and the IP address of the virtualized-information management apparatus are acquired from the virtual volume identifying information illustrated in Fig. 4A. As a result, 20 it is specified which virtualized-information management apparatus 10 manages "volume A".

It is found from the virtual volume sector information in Fig. 4B that the FC address of the storage having the sector corresponding to the virtual volume sector "1 to 1,500" is "021a00". This is because a 25 range of the specified sector is included in the virtual volume sector "1

to 233,444". Further, the real volume sector corresponding to the virtual volume sector "1 to 1,500" is a sector of the storage "1 to 1,500" where the FC address is "021a00". The sector range of the real volume of the storage where the FC address is "021a00" is "1 to 5 233,444", and this corresponds to the sector of the virtual volume of "volume 1".

When the name of the common virtual volume "volume 1" and the virtual volume sector "1 to 1,500" are specified as the access destination of data, the real volume sector of a storage storing the data 10 actually and the virtualized-information management apparatus 10 that manages the real volume of the storage as a virtual volume can be specified. For this reason, the data stored in the real volume sector can be accessed.

As a result of requesting the specified virtualized-information 15 management apparatus 10 for data access, when an access destination changing request is received from the virtualized-information management apparatus 10 due to increase in real volume load, the similar information process is executed on "volume B" where the data of "volume A" are copied. A request is made to the specified 20 virtualized-information management apparatus 10 for data access.

The case of reading data is explained above, but when data are stored into a virtual volume, the similar process is executed so that the virtualized-information management apparatus 10 that manages sectors of corresponding real volumes and a virtual volume configured by the 25 real volumes can be specified. As a result, the data can be stored.

Referring back to Fig. 2, the cache data 232 include data read from the virtual volume by the virtualized-information management apparatus 10. When the virtualized-information management apparatus 10 reads the same data again, it reads the stored data so 5 that the load of the data reading can be reduced.

The control unit 24 totally controls the virtualized-information management apparatus 10, and has a mapping information management unit 240, a real volume management unit 241, a data storage/acquisition unit 242, a data copying unit 243, a virtual volume 10 building unit 244, and a data cache processing unit 245.

The mapping information management unit 240 manages the mapping information 231 in such a manner that it updates the mapping information 231 stored in the storage unit 23 (corresponding information updating unit) based on the updated information of the mapping 15 information 231 received from another virtualized-information management apparatus 10 connected with another SAN 14.

Further, the mapping information management unit 240 transmits/receives the updated information of the mapping information 231 to/from another virtualized-information management apparatus 10 20 connected with another SAN 14 via the IP network 15 (updated information transmission/reception unit). At this time, the updated information of the mapping information 231 is encapsulated with a packet of the IP protocol so as to be transmitted.

When the mapping information management unit 240 receives 25 the updated information from the mapping information 231

encapsulated with the packet of the IP protocol and transmitted from the virtualized-information management apparatus 10, it executes a decapsulating process; so as to fetch the updated information of the mapping information 231.

5        When a new virtual volume is configured, the mapping information management unit 240 adds the information about it to the mapping information 231. When the mapping information management unit 240 accepts a request for deleting a virtual volume, it deletes information about the virtual volume.

10       When data of the virtual volume to be deleted are saved, the data copying unit 243 copies the data into another virtual volume before the virtual volume is deleted.

When the mapping information 231 is updated in such a manner, the updated information is transmitted to another virtualized-information management apparatus 10 connected to another SAN 14, so that the mapping information 231 is shared.

When the request for accessing to the virtual volume is accepted during updating of the mapping information 231, the request is held until the updating is completed, and the access request may be processed after the completion.

The real volume management unit 241 manages the real volume information 230 stored in the storage unit 23. Concretely, when a new real volume is added, the real volume management unit 241 stores the information about the real volume into the real volume information 230, and when the real volume is excluded, it deletes the information about

the real volume.

The data storage/acquisition unit 242 accepts a data storage request or a data acquisition request with respect to a virtual volume from the application server or another virtualized-information management apparatus 10 connected with another SAN 14, and it stores data into a real volume corresponding to the virtual volume or acquires data from the real volume.

5 management apparatus 10 connected with another SAN 14, and it stores data into a real volume corresponding to the virtual volume or acquires data from the real volume.

When the data storage/acquisition unit 242 accepts a request from the application server, it acquires a name of the

10 virtualized-information management apparatus corresponding to the name of the common virtual volume, an IP address of the virtualized-information management apparatus, and a name of the virtual volume in the virtualized-information management apparatus from the virtual volume identifying information illustrated in Fig. 4A

15 based on the information about the name of the specified common virtual volume and the sector of the virtual volume (management apparatus searching unit).

When the specified virtualized-information management apparatus 10 is a self apparatus, the data storage/acquisition unit 242

20 acquires information about the FC address of the storage and the sector of the real volume related with the name of the common virtual volume from the virtual volume sector information illustrated in Fig. 4B.

The data storage/acquisition unit 242 specifies the sector of the real volume mapped in the sector of the specified virtual volume, and stores

25 or acquires data into/from the sector of the real volume.

Further, when the specified virtualized-information management apparatus 10 is another apparatus 10 on another SAN 14, the data storage/acquisition unit 242 transmits the data storage request or the data acquisition request to the virtualized-information management apparatus 10.

When the data storage/acquisition unit 242 accepts the data storage or acquisition request from another virtualized-information management apparatus 10 connected with another SAN 14, it acquires the information about the name of the common virtual volume and the sector of the virtual volume from the virtualized-information management apparatus 10. The data storage/acquisition unit 242 acquires the information about the FC address of the storage and the real volume sector related with the name of the common virtual volume from the virtual volume sector information illustrated in Fig. 4B.

The data storage/acquisition unit 242 specifies the sector of the real volume mapped in the sector of the specified virtual volume, and stores or acquires data into/from the sector of the real volume. When the data acquisition process is executed, the acquired data are transmitted to the virtualized-information management apparatus 10 to which the data acquisition request is transmitted (data storage/acquisition unit).

When data are transmitted to another virtualized-information management apparatus 10 connected with another SAN 14 via the IP network 15, the data storage/acquisition unit 242 encapsulates the data with the packet of the IP protocol so as to transmit them. When the

data storage/acquisition unit 242 receives the data which are encapsulated with the packet of the IP protocol and transmitted from the virtualized-information management apparatus 10 connected with another SAN 14, it executes the decapsulating process so as to fetch 5 the data.

Further, when the data copying unit 243, mentioned later, copies the data stored in the real volume corresponding to the virtual volume into the real volume corresponding to another virtual volume in order to balance the access load on the real volume, the data 10 storage/acquisition unit 242 monitors the load on the real volume corresponding to the virtual volume as a copy source. When the load exceeds a predetermined quantity, the real volume corresponding to the virtual volume as the copy source of the data is specified based on the mapping information 231, and the access destination is switched into 15 the real volume.

The data copying unit 243 copies the data stored in the real volumes configuring the virtual volume into the real volumes configuring another virtual volume (data copying unit). The data may be copied into the real volumes on another SAN.

20 The data are copied by using a known technique for duplicating a volume or a known technique for copying data between the SANs. After the copying, as illustrated in the examples of Figs. 4A and 4B, a request is made to the mapping information management unit 240 for updating the virtual volume identifying information and the virtual 25 volume sector information in the mapping information 231.

A request is made so as to add information about the name of the virtualized-information management apparatus, the IP address of the virtualized-information management apparatus, the name of the virtual volume in the virtualized-information management apparatus, 5 and virtual volume free capacity as the information about the volume in the copying destination of the data to the virtual volume identifying information. Further, a request is made so as to add information about the name of the virtual volume in the virtualized-information management apparatus, the FC address of the storage, and the real 10 volume sector as the information about the volume in the copying destination of the data into the virtual volume sector information.

The virtual volume building unit 244 configures a new virtual volume (virtual storage area configuring unit). When a virtual volume is configured, the virtual volume building unit 244 accepts specifying of 15 condition of properties and a capacity which are satisfied by the virtual volume to be configured. The virtual volume building unit 224 refers to the type of the real volume and the free capacity of the real volume illustrated in Fig. 3B so as to configure the virtual volume which satisfies the conditions.

20 The cache processing unit 245 caches data read from the virtual volume by the virtualized-information management apparatus 10 as the cache data 232 into the storage unit 23. When the same data are read again, the cache processing unit 245 caches the data cached as the cache data 232 to be read from the storage unit 23 (data cache unit).

25 The data to be cached may be data acquired from another

virtualized-information management apparatus 10 connected with another SAN 14 via the IP network 15. As a result, the load of data transmission/reception via the IP network 15 can be reduced.

Fig. 5 is a flowchart of a mapping-information updating procedure by a mapping information management unit 240. When a virtual volume is configured on the SAN 14a connected with the virtualized-information management apparatus 10a, the virtualized-information management apparatus 10a accepts a change of the stored mapping information 231a (step S501). The 10 virtualized-information management apparatus 10a checks whether the mapping information 231a is changed (step S502), and when there is no change (No at step S502), the process is ended.

When the mapping information 231a is changed (Yes at step S502), mapping updating information for updating the mapping 15 information 231a is transmitted to another virtualized-information management apparatus 10b on another SAN 14b (step S503). Since the updated information is transmitted via the IP network, it is encapsulated with the IP protocol so as to be transmitted. As for the updated information to be transmitted, the entire mapping information 20 231a may be transmitted, or only a changed portion may be transmitted.

On the other hand, the virtualized-information management apparatus 10b on another SAN 14b receives the mapping updating information encapsulated with the IP protocol (step S504). The virtualized-information management apparatus 10b fetches the mapping 25 updating information (step S505), and updates the mapping information

231b therein based on the mapping updating information.

When the mapping information 231a in the virtualized-information management apparatus 10a is changed, the mapping updating information is transmitted to the

5 virtualized-information management apparatus 10b on another SAN 14b. In order to check whether the respective virtualized-information management apparatuses 10 on another SANs 14 operate, however, a part of the information stored into the mapping information 231 periodically (for example, the name of the virtualized-information

10 management apparatus or the like) is transmitted/received between the virtualized-information management apparatuses so that the check may be made.

This method is used also when a new SAN and a virtualized-information management apparatus on that SAN are added.

15 Further, when a failure occurs in the IP network 15, the mapping information may be put together in one virtualized-information management apparatus by this method so as to be distributed.

Fig. 6 is a flowchart of a data acquisition procedure by a data storage/acquisition unit 242. The virtualized-information management apparatus 10a accepts a request for accessing to data stored in the virtual volume from the application server 13a or 13b (step S601). The name of common virtual volume and the virtual volume sector illustrated in Fig. 4B are supposed to be specified here.

A check is made based on the name of the specified common

25 virtual volume whether the specified virtual volume is present in one

SAN 14 with reference to the mapping information 231 (step S602).

When the specified virtual volume is present (Yes at step S602), the mapping process for specifying a sector of a real volume into which the data to be accessed are stored is executed (step S603), so that the 5 data are acquired from the sector of the real volume (step S604).

Thereafter, the acquired data are transmitted to the application server 13a or 13b which accepts the data access request (step S605).

When the specified virtual volume is not present in one SAN 14a at step S602 (No at step S602), the virtualized-information management 10 apparatus 10b that manages the virtual volume is searched from the virtual volume identifying information in the mapping information 231 illustrated in Fig. 4A (step S606).

When the virtualized-information management apparatus 10b which manages the virtual volume is specified, the data access request 15 is transmitted to the virtualized-information management apparatus 10b (step S607). In this case, since the data access request is transmitted via the IP network, the data access request is encapsulated with the IP protocol so as to be transmitted.

The virtualized-information management apparatus 10b 20 connected with another SAN 14b waits for the data access request similarly to the virtualized-information management apparatus 10a (step S608). When the virtualized-information management apparatus 10b receives the data access request encapsulated with the IP protocol from the virtualized-information management apparatus 10a (step 25 S609), it deactivates the encapsulation so as to fetch the data access

request (step S610).

Thereafter, the virtualized-information management apparatus 10b executes the mapping process for specifying the sector of the real volume where the data to be accessed are stored with reference to the 5 mapping information 231b in the virtualized-information management apparatus 10b (step S611). The virtualized-information management apparatus 10b acquires data from the sector of the real volume (step S612). When acquiring the data, virtualized-information management apparatus 10b encapsulates the data with the IP protocol so as to 10 transmit them to the virtualized-information management apparatus 10a which has transmitted the data access request (step S613).

The virtualized-information management apparatus 10a receives the data encapsulated and transmitted by the virtualized-information management apparatus 10b (step S614), and deactivates the 15 encapsulation so as to fetch the data (step S615). Thereafter, the sequence proceeds to step S605, and the virtualized-information management apparatus 10a transmits the fetched data to the application server 13a or 13b which has accepted the data access request.

20 Fig. 7 is a flowchart of a virtual-volume building procedure by a virtual volume building unit 244. The virtual volume building unit 244 accepts a virtual volume building request (step S701), and accepts specifying of a condition of the virtual volume to be configured (step S702). The condition of the virtual volume includes a type of a real 25 volume configuring the virtual volume and a necessary capacity of the

virtual volume. The type of the real volume is included as the type of the real volume in the real volume free capacity information illustrated in Fig. 3B.

The virtual volume building unit 244 acquires the real volume information 230 of the SAN 14 where a predetermined priority sequence is the first from the SANs 14 which are used for configuring the virtual volume (step S703). The SAN 14 whose priority sequence is the first is normally the SAN 14 connected with the virtualized-information management apparatus 10 that executes the configuring process.

10 The virtual volume building unit 244 checks whether a real volume which can be used for configuring the virtual volume is present based on the condition accepted at step S702 and the type of the real volume and the real volume free capacity included in the real volume information 230 (step S704).

15 When the usable real volume is present (Yes at step S704), the virtual volume building unit 244 adds the real volume to the virtual volume (step S705), and checks whether the capacity reaches the necessary capacity specified at step S702 (step S706).

When the capacity reaches the necessary capacity (Yes at step S706), the virtual volume building unit 244 updates the mapping information 231 (step S707), and requests the mapping information management unit 240 in the virtualized-information management apparatus 10 connected with another SAN 14 to transmit the mapping updating information (step S708). Thereafter, the virtual volume building unit 244 outputs a notification that configuring of the virtual

volume is completed so as to notify it to an administrator or the like who has made the configuring request.

When no usable real volume is present on the SAN 14 whose priority sequence is the first at step S704 (No at step S704), the virtual 5 volume building unit 244 checks whether the SAN 14 whose priority sequence is the second highest is present (step S710).

When the SAN 14 whose priority sequence is the second highest is present (Yes at step S710), the virtual volume building unit 244 acquires the real volume information 230 about that SAN 14 (step 10 S711). The sequence proceeds to step S704, so that the virtual volume building unit 244 checks whether a usable real volume is present on the SAN 14.

When the SAN 14 whose priority sequence is the second highest is not present (No at step S710), the virtual volume building unit 15 244 outputs a notification that the virtual volume configuration is failed so as to notify it to the administrator or the like who has made the configuration request (step S712).

In the this embodiment, the updated information of the mapping information 231 is transmitted/received between the 20 self-virtualized-information management apparatus 10 and another virtualized-information management apparatus 10 connected with another SAN 14. Further, the mapping information 231 in the self-virtualized-information management apparatus 10 is updated based on the received updated information of the mapping information 231. 25 For this reason, the mapping information 231 which is shared with

another SAN 14 can be held in a dispersed manner. Even if abnormality occurs in the virtualized-information management apparatus 10 on a certain SAN 14, the abnormality does not influence other SANs 14, and thus the virtualizing storage system having high

5 reliability can be configured efficiently.

The embodiment of the present invention is described above, however, variously modified embodiments other than the one described can be made without departing from the scope of the technical spirit of the appended claims.

10 The embodiment explains the FC-SAN where the storage in each SAN transmits/receives the data according to the FC protocol, for example. The present invention is not, however, limited to this, and the invention can be applied to an IP-SAN or the like where the storage in each SAN transmits/receives data according to an Internet small  
15 computer system interface (iSCSI) protocol. Since the iSCSI protocol uses the transmission control protocol/Internet protocol (TCP/IP) for the transmission/reception of data, data can be transmitted/received via the IP network 15 without executing a protocol converting process.

In the case of the IP-SAN, data in virtual volumes can be  
20 accessed by some methods. In a first method, each virtualized-information management apparatus has an initiator function which issues an iSCSI command and a target function which executes the iSCSI command issued by the initiator function.

When the virtualized-information management apparatus  
25 accepts the request for acquiring data in the virtual volume on another

SAN from the application server on the SAN, the apparatus transmits the data acquisition request to the virtualized-information management apparatus on the SAN where the virtual volume is present according to the iSCSI protocol.

5        The virtualized-information management apparatus which receives the data acquisition request acquires data from the real volume specified by executing the mapping process, and transmits the data to the application server via the virtualized-information management apparatus on the same SAN where the application server

10      is present.

In the second method, the application server has the initiator function, and each virtualized-information management apparatus has the target function. When the virtualized-information management apparatus accepts the acquisition request of the data in the virtual volume on another SAN from the application server on the SAN where the apparatus is present, the apparatus transmits the information about the virtual volume where the acquired data are present to the application server.

15      volume on another SAN from the application server on the SAN where the apparatus is present, the apparatus transmits the information about the virtual volume where the acquired data are present to the application server.

20      The application server which receives the virtual volume information transmits the data acquisition request directly to the virtualized-information management apparatus on the SAN where the virtual volume is present according to the iSCSI protocol.

25      The virtualized-information management apparatus which receives the data acquisition request acquires data from the real volume specified by executing the mapping process, and transmits the

data directly to the application server without routing through the virtualized-information management apparatus on the SAN where the application server is present.

In the third method, the application server which receives the  
5 virtual volume information from the virtualized-information management apparatus on the same SAN in the second method accesses directly to the real volume on another SAN based on the information so as to acquire the data. In this case, since the data are written without routing through a cache of the virtualized-information management  
10 apparatus, the cache data 232 should be nullified.

In the first method, the virtualized-information management apparatus on the SAN where the application server which makes the data acquisition request is present is a bottleneck, and thus the access speed is possibly reduced. Since the iSCSI interface, however, does  
15 not have to be provided to the application server, the system can be configured more inexpensively than the second and the third methods.

Meanwhile, in the second and the third methods, data are transmitted/received without routing through the virtualized-information management apparatus on the SAN where the application server is present. For this reason, the volume of communication using the  
20 iSCSI protocol between the virtualized-information management apparatuses can be suppressed, and thus the load on the IP network can be reduced.

In the embodiment, the data such as the updated information of  
25 the mapping information to be transmitted/received are encapsulated

with the IP protocol so as to be tunneled through the IP network. The present invention is, however, not limited to this, and an apparatus having the similar function may be additionally provided.

Fig. 8 is a schematic diagram of a virtualizing storage system 5 when FCIP gateways 16a and 16b execute an encapsulating process. A difference from the virtualizing storage system shown in Fig. 1 is that the FCIP gateways 16a and 16b are provided between the IP network 15 and the SANs 14a and 14b.

The FCIP gateways 16a and 16b encapsulate data of FC 10 protocol with the IP protocol, and decapsulate the encapsulated data of the FC protocol. They can be tunneled through the IP network 15. In this case, it is not necessary that the virtualized-information management apparatuses are present between the IP network 15 and the SANs 14a and 14b.

15 Further, in the embodiment, the virtualized-information management apparatuses connected with the SANs, respectively, have the management function of the mapping information. The present invention, however, is not limited to this, and a domain name system (DNS) server connected with the IP network may have the management 20 function of the mapping information.

In this case, the DNS server manages an IP address of a computer connected with the IP network, and is connected with the IP network, so as to manage the mapping information in the storage where data are transmitted/received by iSCSI.

25 The storage itself may have the management function of the

mapping information. A switch product such as an FC switch may have the management function of the mapping information.

In the embodiment, one virtualized-information management apparatus is connected with one SAN so as to manage the SAN. The 5 present invention, however, is not limited to this, and one virtualized-information management apparatus may manage a plurality of SANs. In this case, the mapping updating information is communicated between the virtualized-information management apparatuses, so that the virtualizing storage system having high 10 reliability can be configured.

As to the respective processes explained in the embodiment, all or some of the processes which are executed automatically can be executed manually, or all or some of the processes which are executed manually can be executed automatically by known methods. The 15 processing procedures, the control procedures, the concrete names, and the information including various data and parameters described in the specification and illustrated in the drawings can be arbitrarily modified unless specially noted.

The components of the respective apparatuses in the drawings 20 are conceptualistic in the functions, and it is not always necessary to physically configure the system as illustrated in the drawings. That is to say, the concrete form of dispersion and integration of the components is not limited to those illustrated in the drawings, and thus all or some of the components can be configured by functionally or 25 physically dispersing and integrating them in an arbitrary unit according

to various loads and use conditions. All or some of the processing functions which are executed in the devices are realized by a central processing unit (CPU) and programs which are analyzed and executed by the CPU, or can be realized as a hardware using a wired logic.

5         The virtualized-information management method explained in the embodiment can be realized by executing a prepared program using a computer such as a personal computer or a workstation. This program can be distributed via the network such as the Internet. The program is recorded into a recording medium readable by the computer 10 such as a hard disk, a flexible disk (FD) a compact-disc read only memory (CD-ROM), a magneto optical disk (MO), and a digital versatile disk (DVD), and the program is read from the recording medium by the computer so as to be executed.

According to the present invention, the information about the 15 updating of the corresponding information is transmitted/received between the self-virtualized-information management apparatus and another virtualized-information management apparatus connected with other network, and the corresponding information of the self-virtualized-information management apparatus is updated based on 20 the received information about the updating of the corresponding information. For this reason, the corresponding information which is shared with another network can be held in a dispersed manner, and even if abnormality occurs in a virtualized-information management apparatus of a certain network, the abnormality does not influence 25 another networks. As a result, the virtualizing storage system having

high reliability can be configured efficiently.

When the data storage request or the data acquisition request with respect to the virtual storage area is accepted from the virtual storage area, a virtualized-information management apparatus which is connected with the network where the virtual storage area is configured is searched. When the searched virtualized-information management apparatus is the self-virtualized-information management apparatus, the data are stored into the real storage area related with the virtual storage area based on the updated corresponding information, or data are acquired from the real storage area related with the virtual storage area. For this reason, the data storage process or the data acquisition process can be executed suitably and efficiently with reference to the updated corresponding information.

When the data acquisition request is accepted from the virtualized-information management apparatus connected with another network, data are acquired from the real storage area related with the virtual storage area based on the updated corresponding information. Further, the acquired data are transmitted to the virtualized-information management apparatus connected with another network. For this reason, the data stored in the storage on the network managed by the virtualized-information management apparatus can be provided to the virtualized-information management apparatus connected with another network efficiently.

When the searched virtualized-information management apparatus is the virtualized-information management apparatus

connected with another network, the data storage request or the data acquisition request is transferred to the searched virtualized-information management apparatus. For this reason, even if the virtual storage area where the data are stored or acquired is not present on the

5 network managed by the self-virtualized-information management apparatus, the request is given to the virtualized-information management apparatus connected with another network. As a result, the data storage or acquisition can be executed efficiently.

When the virtualized-information management apparatus

10 connected with another network receives the transmitted data according to the data acquisition request transferred to the virtualized-information management apparatus, the apparatus stores the data. When the apparatus accepts the data acquisition request next time, it outputs the stored data. For this reason, the load due to the

15 transmission/reception of data can be reduced.

The data stored in the first real storage area is copied into the second real storage area, and the data storage request or the data acquisition request with respect to the first real storage area is accepted. At this time, the data storage or acquisition destination is

20 switched into the second real storage area according to the access load on the first real storage area where the data are stored. For this reason, the access load on the real storage area can be balanced.

The real storage areas are selected based on a predetermined condition, and the selected real storage areas are integrated virtually so

25 as to configure the virtual storage area. Further, the corresponding

information of the self-virtualized-information management apparatus is further updated based on the corresponding information in the configured virtual storage area. For this reason, the virtual storage area having desired properties can be configured efficiently.

5        When the network connected with the self-virtualized-information management apparatus is connected with another network by the IP network, the information about the updating of the corresponding information is encapsulated with the packet of the IP protocol so as to be transmitted. When the information about the 10 updating of the corresponding information encapsulated with the packet of the IP protocol is received, the information about the updating of the corresponding information is decapsulated. For this reason, the system that can transmit/receive the information about the updating of the corresponding information can be configured inexpensively using 15 the widely spread IP network.

When the network connected with the self-virtualized-information management apparatus is connected with another network by the IP network, data are encapsulated with the packet of the IP protocol so as to be transmitted. When the data 20 encapsulated with the packet of the IP protocol are received, the data are decapsulated. Therefore, the system that can transmit/receive data can be configured inexpensively by using the widely spread IP network.

Although the invention has been described with respect to a 25 specific embodiment for a complete and clear disclosure, the appended

claims are not to be thus limited but are to be construed as embodying all modifications and alternative constructions that may occur to one skilled in the art which fairly fall within the basic teaching herein set forth.

5